

The Nature of Dynamical Explanation*

Carlos Zednik^{†‡}

The received view of dynamical explanation is that dynamical cognitive science seeks to provide covering-law explanations of cognitive phenomena. By analyzing three prominent examples of dynamicist research, I show that the received view is misleading: some dynamical explanations are mechanistic explanations and in this way resemble computational and connectionist explanations. Interestingly, these dynamical explanations invoke the mathematical framework of dynamical systems theory to describe mechanisms far more complex and distributed than the ones typically considered by philosophers. Therefore, contemporary dynamicist research reveals the need for a more sophisticated account of mechanistic explanation.

1. Introduction. The theoretical frameworks of computationalism and connectionism are often construed as a search for cognitive mechanisms, the specific structures and processes from which cognitive phenomena arise. In contrast, the framework of dynamicism is generally understood to be a search for principles or laws—mathematical regularities that govern the way cognitive phenomena unfold over time. In recent philosophical discourse, this difference between *traditional* and *dynamical* cognitive science has been framed as a difference in scientific explanation: whereas computationalist and connectionist explanations are *mechanistic explanations*, dynamical explanations take the form of *covering-law explanations* (for discussion, see van Gelder 1995, 1998; Clark 1997, 1998; Bechtel 1998; Bechtel

*Received March 2010; revised August 2010.

†To contact the author, please write to: Indiana University Cognitive Science Program, Eigenmann Building 819, 1910 East 10th Street, Bloomington, IN 47406; e-mail: czednik@indiana.edu.

‡Special thanks for valuable feedback during the preparation of this manuscript go to Adele Abrahamsen, Colin Allen, Randy Beer, Tony Chemero, Paul Williams, and four reviewers at *Philosophy of Science*. This research was supported in part by the National Science Foundation Integrative Graduate Education and Research Traineeship for the Dynamics of Brain-Body-Environment Systems in Behavior and Cognition at Indiana University.

Philosophy of Science, 78 (April 2011) pp. 238–263. 0031-8248/2011/7802-0004\$10.00
Copyright 2011 by the Philosophy of Science Association. All rights reserved.

and Abrahamsen 2002; Chemero and Silberstein 2008; Walmsley 2008; Chemero 2009).

In this article, I challenge the received view of dynamical explanation. After briefly outlining the principles of covering-law explanation and mechanistic explanation in section 2, in section 3 I introduce a well-known dynamical model—the Haken-Kelso-Bunz (HKB) model of coordination dynamics (Haken, Kelso, and Bunz 1985; Kelso 1995)—that is regularly used to articulate and promote the received view. Although this model is often considered to be a paradigmatic example of dynamical explanation, in section 4 I introduce two very different dynamical explanations: Thelen et al.'s (2001) explanation of infant perseverative reaching and Beer's (2003) explanation of perceptual categorization in a *minimally cognitive* agent. Although both of these dynamical explanations invoke the mathematical tools and concepts of dynamical systems theory, they resemble mechanistic explanations. Therefore, the received view is misleading: although the mathematical tools and concepts of dynamical systems theory are frequently used to describe principles and laws, they are also sometimes used to describe cognitive mechanisms.

Insofar as some dynamical explanations are mechanistic in nature, they have much in common with traditional computational and connectionist explanations. Among others, like other mechanistic explanations in cognitive science, these dynamical explanations are reductive explanations (Craver 2007; Bechtel 2008) and may be amenable to *representation hunting* (Chemero and Silberstein 2008). Nevertheless, this does not mean that dynamicist research is not interestingly novel and important. Indeed, in section 5 I show that dynamical explanations are well suited for describing *extended mechanisms* whose components are distributed across brain, body, and the environment. Moreover, I argue that dynamical explanations may be uniquely able to describe mechanisms whose components are engaged in complex relationships of *continuous reciprocal causation* (Clark 1997). Therefore, a closer look at dynamical cognitive science will reveal that the extant philosophical conception of mechanistic explanation may have underestimated practicing scientists' willingness and ability to describe increasingly complex and distributed cognitive mechanisms.

2. Covering-Law Explanation, Mechanistic Explanation, and Traditional Cognitive Science. According to the principles of covering-law explanation, scientific explanation involves the subsumption of a target phenomenon under natural law. That is, a phenomenon is explained when it can be logically deduced from statements of relevant antecedent conditions and one or more laws of nature (Hempel and Oppenheim 1948). Two aspects of this conception are worth highlighting. First, saying that a phenomenon is explained when it can be logically deduced is tantamount

to saying that a phenomenon can be explained when it can be (or could have been) predicted. Second, although “law of nature” has proven to be a notion of significant philosophical controversy, the core idea is one of a highly general, counterfactual-supporting regularity (for discussion, see Cartwright 1983; Salmon 1989). Thus construed, Newton’s explanation of celestial motion is a paradigmatic example of covering-law explanation: his laws of motion and universal gravitation predict not only the Earth’s effect on the motion of the moon but also (purport to) predict the relative effects between *any* physical objects in actual and counterfactual circumstances.

Alas, it is becoming increasingly clear that the principles of covering-law explanation are of limited use for understanding contemporary cognitive science. Although many aspects of cognitive scientific practice are concerned with the discovery of lawlike behavioral regularities, cognitive scientists tend to treat such regularities—sometimes also called *effects*—not as explanations but as explanatory *targets* (Cummins 2000). Indeed, after having been discovered, behavioral regularities are typically explained by describing the mechanisms responsible therefore (Bechtel and Richardson 1993; Machamer, Darden, and Craver 2000; Craver 2007; Wright and Bechtel 2007). One particularly influential conception of *mechanism* is due to William Bechtel: “A mechanism is a structure performing a function in virtue of its component parts, component operations, and their organization. The orchestrated functioning of the mechanism, manifested in patterns of change over time in properties of its parts and operations, is responsible for one or more phenomena” (Bechtel and Abrahamsen 2010, 323; see also Machamer et al. 2000; Bechtel and Abrahamsen 2005; Craver 2007). In this conception, mechanistic explanation consists of *describing* the particular organized collection of parts and operations that is responsible for the behavioral regularity being explained.

Mechanistic explanation is by no means trivial. Although scientists may be able to observe that a complex system exhibits a particular behavioral regularity, they may not know which features of that system are involved in producing it. That is, before it can be described, the mechanism responsible for a behavioral regularity has to be *discovered*. The discovery of mechanisms is frequently facilitated by the explanatory heuristics of *decomposition* and *localization* (Bechtel and Richardson 1993), the former of which comes in two forms. *Structural* decomposition involves breaking a complex system down into a collection of simpler subsystems or parts. But not all structural decompositions lead to the discovery of a mechanism; the parts being revealed should be *working* parts, meaning they should behave in ways that contribute toward the production of the behavioral phenomenon being explained (Craver 2007). Therefore, structural decomposition is usually followed by an attempt to characterize the behavior of individual parts so as to show that they are in fact the working

parts of a mechanism. In contrast, *functional* decomposition involves re-describing the behavioral phenomenon as a series or organized collection of simpler behaviors or operations. Importantly, functional decomposition leads to the discovery of a mechanism only if the operations being posited are actually realized in the system from which the phenomenon arises. One way of showing that this is the case is to localize individual operations in corresponding parts—thus simultaneously showing that the operations are realized in the system and that the parts with which those operations are identified are in fact working parts of a mechanism. Of course, such straightforward localization is not always feasible; scientists frequently have to find more roundabout ways of showing that a given collection of operations is actually realized in the system from which the target phenomenon arises (for discussion, see Bechtel and Richardson 1993; Bechtel and Abrahamsen 2005; Craver 2007).

A better understanding of mechanisms and mechanistic explanation can be attained by considering some examples from traditional cognitive science. In classical computationalism, functional decomposition is achieved through the practice of *computational modeling*, in which the target phenomenon is reproduced by way of an organized series of algorithmic operations on symbolic representations (Cummins 1983). Frequently, computational models are accompanied by a stated commitment, although rarely an explicit attempt, to identify individual operations with specific neurobiological structures—consider David Marr's (1982) attempt to localize the detection of *zero-crossings* in the so-called simple cells of the visual cortex. Therefore, computational models specify the component operations of a mechanism that are then (in ideal cases) localized in neurobiological component parts.

Whereas computational models are used to describe a mechanism's component operations, *artificial neural network models* provide abstract descriptions of the neurobiological systems in which cognitive mechanisms are realized.¹ After a neural network model is trained to reproduce a particular behavioral phenomenon, a variety of mathematical techniques can be used to reveal the mechanism responsible for that phenomenon. Consider Jeffrey Elman's (1991) simple recurrent network model of language processing. After training the network to predict the word or word category that immediately follows an English sentence fragment, Elman

1. Note that the components of an artificial neural network model—nodes and connections—need not correspond to individual neurons and synapses in the brain. Nevertheless, connectionist researchers routinely assume that such models capture basic neurobiological principles such as Hebbian or back-propagation learning, spreading activation, and graceful degradation (Rumelhart et al. 1986; Churchland 1996; Bechtel and Abrahamsen 2002).

conducts a *principal components analysis* of the network's hidden-layer activity to describe three operations: (a) identifying number (singular/plural), (b) identifying noun role (subject/object), and (c) identifying verb type (direct object is required/allowed but not required/precluded). Although each one of these operations is *distributed* across the network's hidden layer and thus cannot be directly localized in any single neural unit, Elman's analysis shows that items *a-c* are in fact the component operations of a mechanism for word prediction and that they are realized in the system from which it arises.²

3. Dynamical Explanation: The Received View. Whereas classical computationalism and connectionism are both typically construed as a search for cognitive mechanisms, dynamicism has traditionally been viewed as a search for principles and laws. Dynamical explanations consist of *dynamical models* (sets of differential or difference equations that capture a particular cognitive system's behavior) and *dynamical analyses* (formal analyses that invoke the tools and concepts of *dynamical systems theory* to describe the modeled system's abstract mathematical properties).³ Consider Scott Kelso's (1995) dynamical explanation of *bimanual coordination*, a robust behavioral phenomenon in which the rhythmic oscillatory motion of two opposing index fingers spontaneously becomes coordinated in a way that depends on the frequency of oscillation. At low frequencies, opposing index fingers reliably settle into one of two patterns of motion: an *in-phase* pattern in which both fingers alternate between simultaneously pointing inward and outward, or an *anti-phase* pattern that resembles the parallel motion of windshield wipers on most cars. At high oscillation frequencies, in contrast, the index fingers settle into the in-phase pattern only.

Kelso's dynamical explanation of bimanual coordination is grounded on the HKB model of coordination dynamics (Haken et al. 1985). This dynamical model consists of a single differential equation that describes in-phase as well as anti-phase motion and that accurately predicts the phase transitions that occur between low and high frequencies:

$$\dot{\varphi} = -a \sin \varphi - 2b \sin 2\varphi.$$

In this model, changes in the between-fingers phase relation φ (where $\varphi = 0^\circ$ corresponds to perfect in-phase motion and $\varphi = 180^\circ$ corresponds to perfect anti-phase motion) are expressed as a function of a and b (where b/a corresponds to the inverse of the index fingers' oscillation frequency).

2. For further discussion of connectionism and mechanistic explanation, see Bechtel (1998), Cummins (2000), Bechtel and Abrahamsen (2002).

3. A formal introduction to dynamical modeling, dynamical analysis, and dynamical systems theory can be found in Strogatz (1994). For a philosophical introduction, see Port and van Gelder (1995).




Figure 1. Vector field diagram of the HKB model with φ as the state variable and b/a as the control parameter. When $b/a > 0.25$, motion is stable when $d\varphi/dt = 0$ and $\varphi = 0^\circ$ (in-phase) or $\varphi = 180^\circ$ (anti-phase). In contrast, when $b/a < 0.25$, motion is stable only when $d\varphi/dt = 0$ and $\varphi = 0^\circ$ (in-phase). An inverse pitchfork bifurcation at $b/a = 0.25$ transforms the anti-phase attractor into an unstable equilibrium point. Figure adapted from Kelso (1995). Color version available as an online enhancement.

A dynamical analysis of the HKB model (fig. 1) shows that bimanual coordination can be understood as a dynamical system with two point attractors when $b/a > 0.25$ (one each corresponding to in-phase and anti-phase motion), one point attractor when $b/a < 0.25$ (in-phase motion), and an inverse pitchfork bifurcation that annihilates the anti-phase attractor when $b/a = 0.25$.

The HKB model accurately describes bimanual coordination and supports a variety of testable predictions. Quantitatively, it predicts the pattern of motion that will be exhibited at any given frequency and initial phase relation. In dynamical terms, for any value of b/a and initial φ , it is possible to determine the attractor basin that will constrain φ 's trajectory through state space. Qualitatively, the model predicts a *hysteresis* effect in which phase transitions occur when the oscillation frequency changes from low values to high but not when the oscillation frequency

changes from high values to low. Once again in dynamical terms, if φ begins in the attractor basin corresponding to anti-phase motion and b/a decreases, it is forced to settle into the in-phase attractor when the anti-phase attractor is annihilated; if, in contrast, φ begins in the attractor basin corresponding to in-phase motion and b/a increases, it remains in that basin indefinitely.

Kelso's dynamical explanation of bimanual coordination is a clear example of covering-law explanation. First, the HKB model is invoked in patterns of deductive inference that conclude with a description of the target phenomenon. Given any particular oscillation frequency and initial phase relation, it can be used to predict the phase relation that will eventually be observed. Second, the HKB model describes a lawlike regularity that is both counterfactual supporting and highly general. On the one hand, this is because the HKB model's predictive power extends to situations in which one or both fingers are forced out of their regular oscillatory motion; although perturbing φ away from its stable in-phase or anti-phase motion temporarily results in a phase relation *between* 0° and 180° , the structure of the attractor landscape predicts that φ will inevitably return to one of those two stable attractors and that it will do so within a predictable amount of time. On the other hand, variations of the same model can be used to describe and predict several kinds of coupled oscillatory motion, including animal locomotion (Kelso 1995), speech production (Port 2003), and behavioral coordination between individuals (Oullier et al. 2005). In other words, the HKB model is not restricted to any particular system but rather describes a "general principle of pattern formation" that applies "regardless of what elements are involved in producing the patterns or at what level these patterns are studied or observed" (Kelso 1995, 2). Thus, in general, Kelso's dynamical explanation of bimanual coordination is an example of covering-law explanation because it relies on a lawlike regularity to deduce properties of the target phenomenon.⁴

But Kelso's dynamical explanation is more than just a well-known example. A brief review of the relevant philosophical literature demonstrates that the features of this particular dynamical explanation are often assumed to be features of dynamical explanation in general. Building on some of the themes in Kelso's (1995) seminal book, Tim van Gelder first introduced the HKB model to the philosophical community "to illustrate the dynamical approach to cognition" (1998, 616) and to draw an analogy between dynamical explanations of cognition and Newton's law-based explanations of celestial motion. Bechtel (1998) was the first to explicitly

4. For further discussion of the HKB model, see Kelso (1995), van Gelder (1998), Walmsley (2008), Chemero (2009).

associate Kelso's brand of dynamicism with the principles of covering-law explanation, and Walmsley recently invoked the HKB model to defend the general claim that "the explanatory goal of dynamical cognitive scientists is to provide covering-law explanations" (2008, 343). Chemero and Silberstein have recently appealed to Oullier et al.'s (2005) adaptation of the HKB model to illustrate the practices of "a growing minority of cognitive scientists" who "have eschewed mechanical explanations and embraced dynamical systems theory" (Chemero and Silberstein 2008, 11), and Chemero has further articulated this view by characterizing the HKB model "as a unifying model and guide to discovery for [dynamical] cognitive science" (2009, 86). These remarks indicate just how profound the influence of Kelso's dynamical explanation has been: It has served not only as a template for future empirical and modeling work (e.g., van Rooij, Bongers, and Haselager 2002) but also as a paradigmatic example on the basis of which general philosophical claims about dynamicism have been articulated and defended. One such claim is the received view of dynamical explanation, the view that dynamical explanation is just a special case of covering-law explanation.

In the next section I will argue that the received view of dynamical explanation is misleading. Kelso's explanation of bimanual coordination is not in fact representative of dynamical explanation in general, and many dynamical explanations actually resemble mechanistic explanations rather than covering-law explanations. In the meantime, however, it is worth noting that the received view renders dynamical explanation susceptible to a well-known philosophical worry. Recall that, according to the principles of covering-law explanation, a phenomenon is explained when it can be predicted from antecedent conditions and one or more lawlike regularities. Frequently, this predictive ability is afforded by lawlike regularities that are *phenomenological* in kind—regularities that are specified primarily in terms of the target phenomenon's observable properties (Cartwright 1983; Craver 2006). The HKB model is a case in point: φ and *b/a* each refer to observable features of bimanual coordination itself, rather than to the underlying neural, biomechanical, or other physical structures and processes from which that phenomenon arises. Now, a well-known philosophical worry about phenomenological laws is that they merely describe rather than genuinely explain. But although this mere description worry is already well established in the theoretical literature on dynamicism (see, e.g., Eliasmith 1996; Dietrich and Markman 2001; van Leeuwen 2005), it is frequently brushed aside by proponents of dynamical explanation who appeal to the fact that even phenomenological laws may support counterfactuals (see, e.g., Clark 1997; Walmsley 2008; Chemero 2009). Because phenomenological laws like the one expressed by the HKB model extend to counterfactual circumstances, they go beyond describing

what is actually observed and therefore do more than merely describe the target phenomenon.

Unfortunately, this simple appeal to counterfactuals fails to address the core intuition behind the mere description worry. In particular, it is easy to think that covering-law explanations based on phenomenological laws leave something important to be desired: “Because phenomenal models summarize the phenomenon to be explained, they typically allow one to answer some [what-if-things-had-been-different] questions. But an explanation shows why the relations are as they are in the phenomenal model, and so reveals conditions under which those relations might change or fail to hold altogether” (Craver 2006, 358). That is, although a phenomenological law may predict what happens in actual and counterfactual circumstances, it often remains unclear why the law applies in the first place. Put differently, phenomenological laws by themselves provide no means of determining when they can or cannot be used in deductive inferences about the target phenomenon. Consider once again Kelso’s dynamical explanation of bimanual coordination. The HKB model can be used to deduce properties of the target phenomenon just because moving index fingers can be understood as a system of *coupled oscillators*. Although this is explicit in Kelso’s (1995) extended discussion of the way in which the HKB model was originally derived, the model itself does not provide any means of determining whether a given system can or cannot be understood as a system of coupled oscillators. In comparison, a (hypothetical) description of the neurobiological and biomechanical mechanisms responsible for bimanual coordination would likely provide a description that *resembles* a system of coupled oscillators. Unlike covering-law explanations based on phenomenological laws, mechanistic explanations provide a means of determining which systems fall under the scope of the explanation and when they do so. As long as dynamical explanation is viewed as a form of covering-law explanation as the received view suggests, the mere description worry looms.

Of course, the mere description worry can be resisted by stipulating that covering-law explanations (even those that are grounded on phenomenological laws) qualify as genuine explanations. This seems to be the view espoused by philosophical proponents of *explanatory pluralism* (see, e.g., Chemero and Silberstein 2008), as well as (implicitly or explicitly) by practicing cognitive scientists who choose to employ various descriptive (explanatory?) techniques to accommodate different kinds of phenomena. Short of adopting this kind of pluralist stance, however, the received view of dynamical explanation implies that dynamicist researchers should more thoroughly engage in well-established philosophical worries about covering-law explanation.

4. Dynamical Explanation as Mechanistic Explanation. Before proceeding to cases of dynamical explanation that are mechanistic rather than covering-law in nature, it is important to recognize that there are no a priori reasons to deny that dynamical models and dynamical analyses can be used to provide mechanistic explanations of cognitive phenomena. As Wright and Bechtel (2007) go to great lengths to impress, mechanistic explanation is an epistemic activity that centers on the practice of describing a mechanism. Such description can be achieved via a variety of descriptive schemes, including verbal characterizations, schematic box-and-arrow diagrams, more-or-less detailed pictures, and physical or simulated two- and three-dimensional models. The differential (or difference) equations that constitute a dynamical model—as well as the mathematical constructs and graphical representations that figure in dynamical analyses—are, I claim, equally viable descriptive schemes. Variables and parameters can be used to describe structural and functional properties (e.g., size, location, velocity, activation) of the parts and operations of a mechanism, and differential equations can be used to capture spatial and temporal relations between those properties as well as to describe the way they change over time.⁵ Likewise, graphical representations such as state-space trajectories, attractor landscapes, bifurcation diagrams, and so on—the representational currency of dynamical analysis—can be used in much the same way. The general point is this: like English, the mathematics of dynamical systems theory is a language, an important feature of which is its capacity to represent. What is being represented—a mechanism, law of nature, or my neighbor's pet iguana—is not determined by the language being used but by the way in which tokens of that language are interpreted. The claim at the heart of my argument is that the differential equations and graphical representations that figure in many dynamical explanations can be, in principle as well as in practice, interpreted as representations of cognitive mechanisms.

One more preliminary issue: there have been at least two other explicit attempts to relate dynamical explanation and mechanistic explanation. Clark (1997) and Bechtel (1998; see also Bechtel and Abrahamsen 2010) each discuss hybrid explanations that combine the mathematical tools and concepts of dynamical systems theory with connectionist and other traditional modeling techniques. Although both Clark and Bechtel understand these hybrid explanations to be mechanistic in nature, what

5. Incidentally, it is also common to describe artificial neural network models as systems of equations (see, e.g., Rumelhart et al. 1986; Elman 1991). Although such models are frequently accompanied by schematic circle-and-arrow diagrams, these diagrams do not (I claim) add significantly to the explanatory power provided by the equations. Rather, they provide alternative (possibly more accessible but typically less detailed) descriptions of the same properties being represented as equations.

renders each one of their examples mechanistic is something other than the mathematics of dynamical systems theory. Consider the explanation of *circadian rhythms* discussed by Bechtel and Abrahamsen (2010). After describing a molecular mechanism that shows how circadian oscillations arise from the negative feedback relationship between *per* mRNA and the PER protein in the nerve cells of *drosophila* (Hardin, Hall, and Rosbash 1990), Bechtel and Abrahamsen introduce a dynamical analysis by Goldbeter (1995) of that mechanism's oscillatory behavior. The purpose of Goldbeter's analysis is to show that the negative feedback relationship produces stable (as opposed to damped) circadian oscillations. That is, it does not itself describe the molecular mechanism but rather analyzes how that mechanism behaves over time. In Bechtel and Abrahamsen's words, dynamical analyses such as Goldbeter's "are not proposals regarding the basic architecture of circadian mechanisms; rather, they are used to better understand the functioning of a mechanism whose parts, operations, and organization already have been independently determined" (2010).

Bechtel and Abrahamsen clearly identify one way in which the dynamical toolkit can be employed in a thoroughly mechanistic cognitive science. Moreover, they recognize that practicing scientists tend to look beyond well-established philosophical dichotomies—in this case, the apparent dichotomy between dynamical and mechanistic explanation—and employ whatever techniques prove most insightful with respect to their explanatory goals. Nevertheless, the hybrid explanation of circadian rhythms does not in fact resemble the dynamical explanations offered by many prominent dynamicist researchers. In what follows, I present two examples in which dynamical models and dynamical analyses are themselves used to describe the parts and operations of a mechanism as well as its organization. Rather than contrasting with mechanistic explanation as the received view suggests, or being complementary to it in the way recently suggested by Bechtel and Abrahamsen, in these examples dynamical explanation is an instance of mechanistic explanation.

As a first example, consider Thelen et al.'s (2001) *dynamical field theory* model of infant perseverative reaching in Jean Piaget's classic A-not-B task:

$$\tau \dot{u}(x, t) = -u(x, t) + S(x, t) + g[u(x'); x].$$

This model describes the way the activation level (u) of every point (x) on a dynamic *motor planning field* changes (over time t) as a function of the field's previous activation ($-u$), an input vector (S), a *cooperativity parameter* (g), and a temporal decay constant (τ). Every point on the field corresponds to a particular spatial location in the A-not-B task environment. If at any moment the activation of a single point increases beyond a particular threshold level, a reach is induced toward the corresponding

location. The likelihood that this threshold level is surpassed depends in part on the value of the cooperativity parameter g , which determines the degree to which any individual point on the field excites or inhibits the activity of its neighbors. Psychologically, g corresponds to a developmental parameter that influences the extent to which maturing infants are able to perform accurate goal-directed reaches at various developmental stages. Appropriately fixing the value of g allows the model to predict performance in perseverative (for infants 7–12 months of age) as well as nonperseverative (12 months and over) episodes of the A-not-B task.

For current purposes, the most significant feature of Thelen et al.'s dynamical model is the fact that the input vector S is a function of three independent inputs: a *task input* that captures the unchanging features of the A-not-B task environment, a *specific input* that corresponds to the cuing event provided during each trial, and a *memory trace* that captures the influence of remembered reaches from earlier trials:

$$S(x, t) = S_{\text{task}}(x, t) + S_{\text{specific}}(x, t) + S_{\text{memory}}(x, t).$$

This tripartite definition of S is significant because it exemplifies the explanatory heuristic of functional decomposition. The individual contributions of task input, specific input, and memory trace can be construed as the (posited) component operations of a mechanism for goal-directed reaching. Expressed as variables, these operations are linked in an equation that captures their role in a mechanism for goal-directed reaching. To be sure, this way of describing a mechanism is fairly abstract—but not any more so than, for example, a flow chart or schematic diagram. Indeed, Machamer et al. (2000; see also Craver 2007) argue that mechanistic explanation frequently starts as a relatively abstract mechanism sketch that leaves ample room for elaboration. In the same spirit, Thelen et al. voice their intention to “speculate further as to possible neuroanatomical areas where such a field might evolve” (2001, 16). Such speculation not only achieves the goal of elaborating on the mechanism sketch provided by their mathematical model but also demonstrates that the authors eventually seek to localize the operations described in the equation in specific neurobiological parts. In other words, Thelen et al. rely on (or at least allude to) the dual heuristics of decomposition and localization to provide a mechanistic explanation of infant perseverative reaching.

Interestingly, Walmsley (2008) articulates two reasons for thinking that Thelen et al.'s dynamical explanation is in fact a covering-law explanation rather than a mechanistic explanation. First, the way the model is expressed—a single mathematical equation—bears a superficial resemblance to many well-known covering-law explanations, including Kelso's dynamical explanation of bimanual coordination. But as I have already suggested above, this is largely inconsequential; what matters is not the

language of description being employed but the way the tokens of that language are interpreted. Second, Walmsley emphasizes the fact that Thelen et al.'s model can be used to derive, by way of deductive inferences, predictions about goal-directed reaching in actual and counterfactual circumstances. But although this is true, it is misleading to emphasize this feature of Thelen et al.'s model over their statement that the model will be used "to explain, . . . *in terms of the normal processes involved in reaching*, behavioral phenomena [in the A-not-B task environment]" (2001, 10; emphasis added). That is, the model is not merely a formalism that can be used to derive predictions about the target phenomenon but is also a means of showing that the complex process of goal-directed reaching arises from the organized activity of a particular collection of simpler processes—specifically, low-level processes of perception and action rather than high-level processes of concept-formation as originally posited by Jean Piaget.

As a second example, consider Randall Beer's (2003) dynamical explanation of perceptual categorization in a simulated brain-body-environment system (fig. 2). The simulated system consists of a single *minimally cognitive* agent, equipped with a 14-neuron continuous-time recurrent neural network (CTRNN) brain and situated in a simple two-dimensional environment that features a single circular or diamond-shaped object. A 16-dimensional dynamical model defines the system's behavior:

$$\tau \dot{s}_i = -s_i + I_i(x, y; \alpha) \quad i = 1, \dots, 7, \quad (1-7)$$

$$\tau \dot{s}_i = -s_i + \sum_{j=1}^7 w_{ji} \sigma[g(s_j + 0)] + \sum_{j=8}^{12} w_{ji} \sigma(s_j + 0_j) \quad i = 8, \dots, 12, \quad (8-12)$$

$$\tau \dot{s}_i = -s_i + \sum_{j=8}^{12} w_{ji} \sigma(s_j + \theta) \quad i = 13, 14, \quad (13, 14)$$

$$\dot{x} = 5 [\sigma(s_{13} + \theta_{13}) - \sigma(s_{14} + \theta_{14})], \quad (15)$$

$$\dot{y} = -3. \quad (16)$$

Equations (1)–(16) define the change over time in the brain's neural activity ($s_1 \dots s_{14}$), the agent's horizontal position (x), and the object's vertical position (y). Neural parameters w , τ , σ , and θ are predetermined by an artificial evolutionary process and remain constant during as well as between trials. In contrast, the brain's neural activity is continuously affected by the changing sensory input vector I , a function of shape parameter α and of the relative positions of agent and object. Over the




Figure 2. The simulated agent, (A) its continuous-time neural network brain, and (B) its task environment. Figures adapted from Beer (2003).

course of a single trial, a circular or diamond-shaped object falls vertically toward the agent, to which the agent responds by moving horizontally to catch circles and avoid diamonds, thus performing a categorical discrimination. Notably, discrimination is always preceded by an episode of *active scanning*: as the object falls, the agent repeatedly moves from side to side before eventually settling on a position either directly beneath the object or away to one side. Because this active-scanning behavior is an unexpected result of the artificial evolutionary process and cannot be straightforwardly read off the equations of the dynamical model, it constitutes a target phenomenon of significant explanatory interest.

Beer's goal is to use the tools and concepts of dynamical systems theory to explain how the observed active-scanning behavior (and thus, perceptual categorization) arises from the brain-body-environment system defined by the equations of the dynamical model. To this end, he adopts a clearly stated decompositional strategy: "We will decompose the agent-environment dynamics into: (1) the effect that the relative positions of the object and the agent have on the agent's motion; (2) the effect that the agent's motion has on the relative positions of the object and the agent" (Beer 2003, 228; see also Beer 1995).

This brief statement requires elaboration. In dynamical terms, equations (1)–(16) define a dynamical system composed of the simulated agent, its CTRNN brain, and the falling object. This dynamical system is mathematically equivalent to a pair of *coupled* dynamical systems (henceforth labeled "*B*" and "*E*") defined by equations (1)–(14) and (15)–(16), respectively. Whereas *B* is a model of the embodied CTRNN brain that transforms sensory input into motor output, *E* is a model of the environment in which motor output is converted into sensory input. The couplings between *B* and *E* are defined by variables s_{13} and s_{14} in equation (14) and x and y in equations (1)–(7) and are best understood as a model of the two-way interface—the agent's body—that mediates between the brain and the environment. On this interpretation, Beer's aim is to show how the phenomenon of active scanning arises from the simultaneous and coupled activity of the embodied brain on the one hand and of the environment on the other.

At the heart of Beer's dynamical explanation lies a dynamical analysis that describes the activity of each of the two components—embodied brain and environment—identified above. This analysis comes in the form of a pair of steady state velocity fields and superimposed *motion trajectories* (fig. 3). The steady state velocity fields (*colored regions*) describe *E*'s effect on the qualitative behavior of *B*—how the sensory input received at a particular position (x, y) constrains \dot{x} , the agent's horizontal velocity. The motion trajectories (*colored lines*) describe *B*'s effect on the state of *E*—how \dot{x} causes a change in x and thus a change in the environment defined by (x, y) .

The motion trajectories are superimposed on the steady state velocity fields in order to describe the spatiotemporal organization of embodied brain and environment. Of particular importance is the way the motion trajectories *overshoot* some colored regions while performing *reversals* over others. What determines whether a particular motion trajectory performs an overshoot or a reversal is the color of the region over which it is moving, as well as the amount of time it spends in that region. Specifically, a motion trajectory of a particular color (*red or blue*) performs a reversal whenever it is situated over a region of the opposite color and remains




Figure 3. Steady state velocity fields with superimposed motion trajectories for circle trials (*left*; “catch”) and diamond trials (*right*; “avoid”). Axes designate the relative positions of agent (x) and object (y), which together define the state of the environment. Colored regions describe the effect of the environment on the qualitative behavior (steady state velocity) of the embodied brain: blue indicates a steady state velocity directed toward the object; red indicates a steady state velocity directed away from the object; green designates a steady state velocity that may go in either direction. Colored lines correspond to a sampling of recorded motion trajectories several starting locations. Their color indicates the agent’s instantaneous horizontal velocity according to the same color scheme as above, their shape indicates the way the relative positions of object and agent change over time, thus determining the agent’s sensory input. Figure courtesy of Beer (2003). Color version available as an online enhancement.

in it long enough for the agent’s instantaneous horizontal velocity to approach the steady state velocity indicated by that region’s color (red indicates velocities directed away from the object; blue indicates velocities directed toward the object).⁶ Because active scanning is nothing but a particular combination of overshoots and reversals, it is explained by the particular details—shape and color—of the motion trajectories and steady state velocity fields in figure 3. That is, perceptual categorization via active scanning arises from the detailed ways in which B converts sensory input into motor output and E allows motor output to feed back on sensory input.

Although this discussion provides merely a rough sketch of Beer’s dy-

6. The exact amount of time required for this to happen is a function of the neural decay constant τ . If τ is high, the agent’s instantaneous velocity will approach the prescribed steady state velocity quickly and, therefore, quickly perform a reversal when trajectory and region are of opposite color (see Beer [2003] for details).

namical explanation of perceptual categorization via active scanning (but see sec. 5.2 below), it suffices to show that the explanation is mechanistic in nature. First, Beer relies on the explanatory heuristic of structural decomposition to identify two working parts—the embodied brain on the one hand and the environment on the other—of a two-component mechanism realized in the simulated brain-body-environment system. Then, he provides a detailed dynamical analysis to describe the operations associated with each part—the way the embodied brain’s steady state constrains motion in the environment and the way that motion affects the embodied brain’s steady state at every possible relative position of agent and object. Finally, Beer describes the mechanism’s spatiotemporal organization by discussing the detailed relationships between steady state velocity fields and motion trajectories that lead to overshoots and reversals. In summary, therefore, Beer describes the component parts, component operations, and organization of a mechanism for perceptual categorization via active scanning.⁷

One additional feature of Beer’s dynamical explanation is worth highlighting. Although the mechanistic explanation of perceptual categorization via active scanning is contained in his description of the mechanism composed of *B* and *E*, the behavior of at least one of those components affords further explanation. Why does the embodied brain behave the way it does in response to changes in the environment? Beer answers this question by describing the contributions of individual neurons in the agent’s CTRNN brain. To this end, he provides separate steady state velocity fields for interneurons s_8 , s_9 , s_{11} , and s_{12} (fig. 4) and shows that their sum is approximately equivalent to the steady state velocity field on the left side of figure 3.⁸ That is, he describes a *neural-level* mechanism to explain the behavior of one of the two components of the agent-level mechanism responsible for active scanning.⁹ In this sense, Beer’s expla-

7. Note that the plausibility of this construal depends on two critical assumptions I will defend in sec. 5 below: that the components of a single mechanism may span the boundaries between brain, body, and the environment (sec. 5.1) and that the operations associated with the components of a mechanism may take the form of continuous reciprocal interactions (sec. 5.2).

8. The equivalence is approximate because the agent’s motor output is not a linear combination of the four interneuron outputs (Beer 2003, 233). Note that interneuron s_{10} is omitted from this analysis because it does not appear to “play any significant role in the circuit’s operation” (231).

9. Importantly, Beer’s explanation of active scanning is not mechanistic *because* it describes a neural-level mechanism. The description of the neural-level mechanism does not explain the phenomenon of perceptual categorization via active scanning but rather explains the behavior of *B* as described in the steady state velocity fields of fig. 3 (for further discussion, see sec. 5.1).




Figure 4. Steady state velocity fields for interneurons 1–4 of the agent’s CTRNN “brain.” Axes and coloring schemes as in figure 3. The color at any (x, y) in each of the four fields indicates the corresponding interneuron’s contribution to the agent’s steady state horizontal velocity during circle (“catch”) trials. Summing these four interneuron steady state velocity fields yields a steady state velocity field that closely resembles the one on the left side of figure 3. Figure courtesy of Beer (2003). Color version available as an online enhancement.

nation is hierarchical in a way that closely resembles mechanistic explanations in other scientific domains; the components of a mechanism at one level of organization are further decomposed and explained in terms of mechanisms at lower levels of organization (Machamer et al. 2000; Craver 2007; Bechtel 2008).

It is time to take stock. Thelen et al. (2001) and Beer (2003) each offer a dynamical explanation of a (minimally) cognitive phenomenon. In each case, the explanation proceeds by identifying the component parts and operations of a mechanism and by showing how the organized activity of these parts and operations produces the phenomenon being explained. Therefore, Thelen et al. and Beer each provide a counterexample to the received view of dynamical explanation: some dynamical explanations are mechanistic explanations rather than covering-law explanations.

5. Cognitive Mechanisms: Extended and Complex? Section 4 shows that dynamical cognitive science sometimes seeks the description of cognitive mechanisms. In this sense, many branches of dynamical cognitive science resemble computationalist and connectionist cognitive science as discussed in section 2. Although this may appear disappointingly conservative to the more radical proponents of dynamicism, I do not intend to deny that dynamicism is interestingly novel and important. Indeed, I agree with the popular assessment that dynamicist researchers are far more willing and able than some of their traditional colleagues to account for cognitive phenomena that arise from complex reciprocal interactions between brain, body, and the environment (for discussion, see Thelen and Smith 1994; Beer 1995, 2003; Kelso 1995; van Gelder 1995, 1998; van Gelder and Port

1995; Clark 1997; Chemero and Silberstein 2008; Chemero 2009). Nevertheless, in this section I argue that this willingness and ability does not follow from a rejection of mechanistic explanation but from the invocation of the mathematical tools and concepts of dynamical systems theory. Indeed, a closer look at the dynamical explanations introduced in section 4 will show that the extant philosophical conception of mechanistic explanation may be underestimating the degree to which cognitive scientists are able to describe increasingly complex and distributed cognitive mechanisms.

5.1. Extended Mechanisms. To a large extent, the philosophical conception of mechanistic explanation in cognitive science has been articulated with a view to the work of cognitive neuroscientists who seek to describe brain-bound neurobiological mechanisms (see, e.g., Machamer et al. 2000; Craver 2007; Bechtel 2008). Although it is widely recognized that mechanistic explanation also occurs at more abstract levels of analysis (e.g., those relevant to the computational and artificial neural network models discussed in sec. 2), it is a widespread assumption that cognitive mechanisms are localized entirely within biological brains. Alas, this assumption is at odds with the claim that Beer's dynamical explanation is mechanistic in nature. Recall that although *B*—the embodied brain—is responsible for converting sensory input into motor output, the phenomenon of active scanning arises only from *B*'s interaction with *E*, the environment through which motor output feeds back on sensory input. Insofar as it makes sense to talk of *B* and *E* as the components of a (minimally) cognitive mechanism, that mechanism is *extended*; its components are distributed across brain, body, and environment.

Although previously unarticulated, the idea that cognitive mechanisms may be extended is consistent with the extant philosophical conception of mechanisms and mechanistic explanation. Carl Craver (2007, 2009) has recently explored the question of how scientists determine the boundaries of mechanisms and has argued that they do so not via considerations of spatial proximity or between-component interactivity but “by reference to the phenomenon that the mechanism explains” (Craver 2007, 123). More precisely: “Within the boundaries of a mechanism are all and only the entities, activities, and organizational features relevant to the phenomenon selected as our explanatory, predictive, or instrumental focus” (Craver 2009, 590–91). Therefore, insofar as the phenomenon of interest is *perceptual categorization via active scanning* (as opposed to, e.g., *the conversion of sensory input into motor output*), and insofar as active scanning results from the interactions between the embodied brain and the environment as suggested by Beer, the mechanism for active scanning must span the boundaries between simulated brain, body, and environ-

ment. To further illustrate this point, consider the fact that although the steady state velocity fields in figure 3 can be reconstructed from the interneuron steady state velocity fields in figure 4, the latter only describe one of the two operations that contribute to the observed behavior; active scanning can only be reconstructed by superimposing the steady stated velocity fields with motion trajectories. Therefore, a description of the environment is essential to the explanation of the target phenomenon.

Beer's extended mechanism is unlikely to be an isolated example. Many empirical findings in the literature on "embodied and embedded cognition" suggest that extended mechanisms figure far more prominently in the explanatory practices of cognitive scientists than the extant philosophical conception of mechanistic explanation acknowledges. Andy Clark (2008) has recently suggested that cognitive scientists frequently rely on the explanatory heuristic of *distributed* functional decomposition—roughly, the heuristic of functional decomposition outlined in section 2 applied to the behavior of systems that span the physical boundaries between brain, body, and the environment. Rumelhart et al.'s (1986) well-known reflections on chalkboard-based long division, Kirsh and Maglio's (1994) studies of Tetris game-playing, and Thelen and Smith's (1994) explanation of infant treadmill-stepping, among others, can each be thought to employ this heuristic. Although many of these studies have yet to be worked out in as much detail as the canonical examples of mechanistic explanation, such as the explanations of the action potential (Machamer et al. 2000; Craver 2007) and the Krebs cycle (Bechtel 1998, 2008), it is no stretch to think that they too seek to describe "structure[s] performing a function in virtue of [their] component parts, component operations, and their organization" (Bechtel and Abrahamsen 2010, 323)—albeit ones whose parts, operations, and organizations are distributed across physical boundaries. Therefore, I submit that an adequate account of scientific explanation in embodied and embedded cognitive science will have to take seriously the notion of extended mechanisms.

To what extent have philosophers of science already considered the possibility that cognitive mechanisms may be extended? Bechtel (2009) acknowledges that the environment plays an important role in providing the background conditions necessary for mechanisms to perform their adaptive function and that mechanisms are frequently responsible for governing an organism's interaction with its environment. Nevertheless, he denies that cognitive mechanisms themselves extend into the environment: "For *mental* phenomena it is appropriate to treat the mind/brain as the locus of the responsible mechanism and to emphasize the boundary between the mind/brain and the rest of the body and between the cognitive agent and its environment" (Bechtel 2009, 156).

Bechtel's emphasis on "mental" is meant to highlight "behavioral or

psychological” phenomena, as opposed to the “prototypically social” ones in which “the agent is so intertwined with entities outside itself that the responsible system includes one or more cognitive agents and their environment” (Bechtel 2009, 156). But if the empirical findings from the literature on embodied and embedded cognition are to be taken seriously, Bechtel is wrong to draw such a strong distinction between the social and the mental; both kinds of phenomena can arise from systems in which organisms are strongly intertwined with their environments. Accordingly, extended mechanisms are just as likely to figure in explanations of embodied and embedded cognition as they are in explanations of social cooperation. Philosophers wishing to understand the nature of these explanations are well advised to take note.

Finally, it is worth considering the potential importance of dynamical modeling and dynamical analysis to the description of extended mechanisms. As others have already suggested, “dynamical systems theory is especially appropriate for explaining cognition as interaction with the environment because single dynamical systems can have parameters on each side of the skin” (Chemero and Silberstein 2008, 14; see also Clark 1997). I see no reason to assume that this appropriateness—afforded primarily by the abstract nature of the dynamical toolkit—applies any differently to the description of extended mechanisms than it does to the characterization of principles or laws. As Beer demonstrates, it is possible to define equations to describe the changing state of the brain and couple them to equations that capture the changing state of the body or the environment. Indeed, by using a single mathematical language to describe brain, body, and environment, cognitive scientists will be able not just to develop a formal understanding of individual components of extended mechanisms but additionally to develop a formal understanding of the ways in which such components interact.

5.2. Dealing with Continuous Reciprocal Causation. There is a second way in which a closer look at contemporary dynamicist research should influence the philosophical conception of mechanistic explanation. Recall that the dynamical system defined by equations (1)–(16) is equivalent to a pair of *coupled* dynamical systems, *B* and *E*, defined by equations (1)–(14) and (15)–(16), respectively. *Coupling* is a technical term that applies whenever two or more dynamical systems mutually influence one another’s change over time. In the philosophical literature, such mutual influence is more commonly known as *continuous reciprocal causation* (Clark 1997). Systems *B* and *E* are engaged in a relationship of continuous reciprocal causation because each system’s behavior is at all times determining, as well as being determined by, the other’s. Despite the presence of continuous reciprocal causation, however, Beer’s dynamical analysis

adequately describes the mechanism for perceptual categorization via active scanning.

The significance of this feat is not to be missed; continuous reciprocal causation is commonly thought to *preclude* mechanistic explanation (see, e.g., van Gelder 1995; Clark 1997; Wheeler 2005). In particular, the presence of such a relationship is thought to impose limits on the explanatory heuristics of structural and functional decomposition: “With increasing continuous reciprocal causation, it will become progressively more difficult both to specify distinct and robust causal-functional roles played by reliably identifiable parts of the system, and to explain interesting system-level behavior in terms of the properties of a small number of subsystems” (Wheeler 2005, 260–61).

Whether or not it is possible to divide a system into a collection of structural parts, the presence of continuous reciprocal interactions is likely to prevent researchers from describing how any one of those parts contributes to the system’s behavior as a whole—it will prevent them from showing that parts are in fact working parts. Similarly, although it may be possible to analyze a complex system’s behavior into a set of functional operations, continuous reciprocal interactions make it difficult or impossible to allocate responsibility for any particular operation to one part of the system, as opposed to allocating responsibility to the system as a whole. Because continuous reciprocal causation seems to impose limits on the heuristics of decomposition and localization, it is frequently thought to preclude mechanistic explanation.

Although it is by no means original to propose dynamical explanation as a way of dealing with continuous reciprocal causation (see, e.g., Thelen and Smith 1994; Kelso 1995; van Gelder 1995; Clark 1997; van Orden, Holden, and Turvey 2003; Wheeler 2005), I suggest that dynamical systems theory can be used to deal with this relationship from within the framework of mechanistic explanation. The reason dynamical systems theory can be used in this way is that its tools and concepts are purpose-built to describe the qualitative behavior of a system in a way that does not depend on a prior description in precise quantitative terms. Once again, Beer’s (minimally) cognitive mechanism for active scanning can be used to illustrate this distinction. Because *B* and *E* are engaged in a relationship of continuous reciprocal causation, it is effectively impossible to make precise quantitative predictions about either system’s behavior over time. Nevertheless, the steady state velocity fields in figure 3 describe an important qualitative feature of *B* relative to the possible states of *E*—the velocity the agent would assume if (counterfactually) its sensory input were to be held constant for an extended period of time.

Does this kind of description suffice for the purposes of mechanistic explanation? Craver suggests that an adequate description of a mechanism

should render it “potentially exploitable for the purposes of manipulation and control” (2007, 94). That figure 3 satisfies this requirement is evidenced by the fact that Beer explains the difference between catching and avoiding by referring to specific details in the steady state velocity fields. Manipulating the model system in a way that changes the size of certain colored regions leads to (predictively) novel behavior: while decreasing the size of the central black region leads to more frequent (correct as well as incorrect) avoiding, increasing it at the expense of the size of the surrounding red regions leads to more frequent catching (Beer 2003, 228–30). Moreover, although not explicitly mentioned by Beer, similar changes in performance can be achieved by manipulating the velocity of the agent’s horizontal motion; increasing the velocity would inevitably lead to more frequent overshoots of individual regions as well as of the midline of figure 3, thus affecting the agent’s catch/avoid behavior. Because figure 3 allows for this kind of predictive reasoning, it provides an adequate description of the operations of the mechanism for perceptual categorization via active scanning.

The moral of the story is that the tools and concepts of dynamical systems theory can be used to describe mechanisms that exhibit continuous reciprocal causation. Although important questions do remain about the degree to which Beer’s methods will scale up to larger and increasingly realistic systems in which continuous reciprocal causation is increasingly prevalent, Beer’s analysis shows that continuous reciprocal causation does not necessarily preclude mechanistic explanation. Apparently, while philosophers have spent their time worrying about the threat of continuous reciprocal causation, practicing dynamicist researchers have busied themselves developing ways to meet the threat head-on.

6. Conclusion. Dynamical explanation is undeniably novel and important. As has already been argued elsewhere (e.g., Elman 1991; Thelen and Smith 1994; van Gelder and Port 1995; Bechtel and Abrahamsen 2010), in addition to introducing a very different mathematical framework to the study of cognitive phenomena, dynamical explanations are able to offer new perspectives on the role of temporal structure in behavior and cognition. In section 5, I argued that dynamical explanations are also uniquely well suited for dealing with continuous reciprocal causation and with interactions between brain, body, and the environment. As it happens, although in some well-known cases these explanations specify principles or laws, in other cases they describe the component parts, operations, and organization of mechanisms. Therefore, contrary to the received view, some dynamical explanations are mechanistic explanations rather than covering-law explanations. Moreover, contrary to the extant philosophical conception of mechanistic explanation, cognitive science may

after all be in a position to describe mechanisms that exhibit continuous reciprocal causation and whose components span the boundaries between brain, body, and the environment.

In closing, it is worth noting some of the broader conceptual ramifications of the claim that some dynamical explanations are mechanistic in nature. First, mechanistic explanation is a form of reductive explanation; phenomena manifested at one level of organization are explained in terms of component parts and operations at lower level(s) of organization (Machamer et al. 2000; Craver 2007; Bechtel 2008). Because some dynamical explanations are mechanistic explanations, dynamical cognitive science is not as thoroughly antireductionist as many early proponents of dynamicism have claimed (see, e.g., Thelen and Smith 1994; Kelso 1995; van Gelder 1998). Rather than being solely interested in identifying the principles that govern behavior and cognition at the “highest relevant level of causal organization” (van Gelder 1998, 622), dynamicist researchers seem equally interested in questions concerning the ways in which such regularities result from structures and processes at lower levels of organization.

Second, those dynamicist researchers who seek to provide mechanistic explanations rather than covering-law explanations may be steering toward reconciliation with proponents of *representationalism*. By describing cognitive mechanisms rather than principles or laws, these researchers describe structures that are amenable to what Chemero and Silberstein (2008) have called *representation hunting*—characterizing the components of a mechanism as representation producers and representation consumers and understanding their operations in terms of the transfer and manipulation of information. Notably, episodes of representation hunting have already occurred both for Beer’s dynamical agent (Ward and Ward 2009) and for Thelen et al.’s dynamical field theory model (Spencer and Schönner 2003). Although true reconciliation between dynamicism and representationalism will be possible only if a suitable notion of *representation* can be articulated—one that accommodates the emphasis on temporal structure, continuity, and physical distribution distinctive of dynamical cognitive science—the search for such a notion is very likely to be intriguing.

REFERENCES

- Bechtel, W. 1998. “Representations and Cognitive Explanations: Assessing the Dynamicist Challenge in Cognitive Science.” *Cognitive Science* 22 (3): 295–317.
- . 2008. *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. London: Routledge.
- . 2009. “Explanation: Mechanism, Modularity, and Situated Cognition.” In *Cambridge Handbook of Situated Cognition*, ed. P. Robbins and M. Aydede, 155–70. Cambridge, MA: Cambridge University Press.
- Bechtel, W., and A. Abrahamsen. 2002. *Connectionism and the Mind: Parallel Processing, Dynamics, and Evolution in Networks*. 2nd ed. Oxford: Blackwell.

- . 2005. "Explanation: A Mechanist Alternative." *Studies in History and Philosophy of Biological and Biomedical Sciences* 36:421–41.
- . 2010. "Dynamic Mechanistic Explanation: Computational Modeling of Circadian Rhythms as an Exemplar for Cognitive Science." *Studies in History and Philosophy of Science A*, 1:321–33.
- Bechtel, W., and R. C. Richardson. 1993. *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton, NJ: Princeton University Press.
- Beer, R. D. 1995. "Computational and Dynamical Languages for Autonomous Agents." In Port and van Gelder 1995, 121–47.
- . 2003. "The Dynamics of Active Categorical Perception in an Evolved Model Agent." *Adaptive Behavior* 11 (4): 209–43.
- Cartwright, N. 1983. *How the Laws of Physics Lie*. New York: Clarendon.
- Chemero, A. 2009. *Radical Embodied Cognitive Science*. Cambridge, MA: MIT Press.
- Chemero, A., and M. Silberstein. 2008. "After the Philosophy of Mind: Replacing Scholasticism with Science." *Philosophy of Science* 75:1–27.
- Churchland, P. M. 1996. *The Engine of Reason, the Seat of the Soul*. Cambridge, MA: MIT Press.
- Clark, A. 1997. *Being There: Putting Brain, Body, and World Back Together Again*. Cambridge, MA: MIT Press.
- . 2008. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: Oxford University Press.
- Craver, C. 2006. "When Mechanistic Models Explain." *Synthese* 153:355–76.
- . 2007. *Explaining the Brain*. Oxford: Oxford University Press.
- . 2009. "Mechanisms and Natural Kinds." *Philosophical Psychology* 22 (5): 575–94.
- Cummins, R. 1983. *The Nature of Psychological Explanation*. Cambridge, MA: MIT Press.
- . 2000. "'How Does It Work?' versus 'What Are the Laws?': Two Conceptions of Psychological Explanation." In *Explanation and Cognition*, ed. F. Keil and R. Wilson, 117–44. Cambridge, MA: MIT Press.
- Dietrich, E., and A. Markman. 2001. "Dynamical Description versus Dynamical Modeling: Reply to Chemero." *Trends in Cognitive Sciences* 5 (8): 332.
- Eliasmith, C. 1996. "The Third Contender: A Critical Examination of the Dynamicist Theory of Cognition." *Philosophical Psychology* 9 (4): 441–63.
- Elman, J. 1991. "Distributed Representations, Simple Recurrent Networks, and Grammatical Structure." *Machine Learning* 7:195–224.
- Goldbeter, A. 1995. "A Model for Circadian Oscillations in the *Drosophila* Period Protein (PER)." *Proceedings of the Royal Society of London B* 261 (1362): 319–24.
- Haken, H., J. A. S. Kelso, and H. Bunz. 1985. "A Theoretical Model of Phase Transitions in Human Hand Movements." *Biological Cybernetics* 51:347–442.
- Hardin, P. E., J. C. Hall, and M. Rosbash. 1990. "Feedback of the *Drosophila* Period Gene Product on Circadian Cycling of Its Messenger RNA Levels." *Nature* 343 (6258): 536–40.
- Hempel, C., and P. Oppenheim. 1948. "Studies in the Logic of Explanation." *Philosophy of Science* 15 (2): 135–75.
- Kelso, J. A. S. 1995. *Dynamic Patterns: The Self-Organization of Brain and Behavior*. Cambridge, MA: MIT Press.
- Kirsh, D., and P. Maglio. 1994. "On Distinguishing Epistemic from Pragmatic Action." *Cognitive Science* 18:513–49.
- Machamer, P., L. Darden, and C. Craver. 2000. "Thinking about Mechanisms." *Philosophy of Science* 67:1–25.
- Marr, D. 1982. *Vision*. New York: Holt.
- Oullier, O., G. C. de Guzman, K. J. Jantzen, J. F. Lagarde, and J. A. S. Kelso. 2005. "Spontaneous Interpersonal Synchronization." In *European Workshop on Movement Sciences: Mechanics-Physiology-Psychology*, ed. C. Peham, W. I. Schöllhorn, and W. Verwey, 34–35. Köln: Sportverlag.
- Port, R. 2003. "Meter and Speech." *Journal of Phonetics* 31:599–611.
- Port, R., and T. J. van Gelder. 1995. *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, MA: MIT Press.

- Rumelhart, D. E., P. Smolensky, J. L. McClelland, and G. E. Hinton. 1986. "Schemata and Sequential Thought Processes in PDP Models." In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, vol. 2, *Psychological and Biological Models*, ed. D. E. Rumelhart, J. E. McClelland, and the PDP Research Group, 7–57. Cambridge, MA: MIT Press.
- Salmon, W. 1989. *Four Decades of Scientific Explanation*. Minneapolis: University of Minnesota Press.
- Spencer, J. P., and G. Schöner. 2003. "Bridging the Representational Gap in the Dynamic Systems Approach to Development." *Developmental Science* 6:392–412.
- Strogatz, S. 1994. *Nonlinear Dynamics and Chaos*. Cambridge, MA: Westview.
- Thelen, E., G. Schöner, C. Scheier, and L. B. Smith. 2001. "The Dynamics of Embodiment: A Field Theory of Infant Perservative Reaching." *Behavioral and Brain Sciences* 24:1–34.
- Thelen, E., and L. B. Smith. 1994. *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: MIT Press.
- van Gelder, T. J. 1995. "What Might Cognition Be, If Not Computation?" *Journal of Philosophy* 91 (7): 345–81.
- . 1998. "The Dynamical Hypothesis in Cognitive Science." *Behavioral and Brain Sciences* 21:1–14.
- van Gelder, T. J., and R. Port. 1995. "It's about Time: An Overview of the Dynamical Approach to Cognition." In Port and van Gelder 1995, 1–44.
- van Leeuwen, M. 2005. "Questions for the Dynamicist." *Minds and Machines* 15:271–333.
- van Orden, G. C., J. G. Holden, and M. T. Turvey. 2003. "Self-Organization of Cognitive Performance." *Journal of Experimental Psychology: General* 132:331–50.
- van Rooij, I., R. Bongers, and W. Haselager. 2002. "A Non-representational Approach to Imagined Action." *Cognitive Science* 26:345–75.
- Walmsley, J. 2008. "Explanation in Dynamical Cognitive Science." *Minds and Machines* 18 (3): 331–48.
- Ward, R., and R. Ward. 2009. "Representation in Dynamical Agents." *Neural Networks* 22:258–66.
- Wheeler, M. 2005. *Reconstructing the Cognitive World*. Cambridge, MA: MIT Press.
- Wright, C., and W. Bechtel. 2007. "Mechanisms and Psychological Explanation." In *Philosophy of Psychology and Cognitive Science*, ed. P. Thagard, 31–79. New York: Elsevier.